End-to-end vs. Hop-by-hop Transport

Simon Heimlicher
Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland
heimlicher@tik.ee.ethz.ch

Pavan Nuggehalli
Centre for Electronics Design and Technology
Indian Institute of Science, Bangalore, India
pavan@cedt.iisc.ernet.in

Martin May
Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland
may@tik.ee.ethz.ch

The transport layer has been considered an end-to-end issue since the early days of the Internet in the 1980s [1], when the TCP/IP protocol suite was designed to connect networks of dedicated routers over wired links. However, over the last quarter of a century, network technology as well as the understanding of the Internet has changed, and today's wireless networks differ from the Internet in many aspects. Since wireless links are unreliable, it is often impossible to sustain an end-to-end connection to transmit data in wireless network scenarios. Even if an end-to-end path exists in the network topology for some fraction of the communication, it is likely to break due to singal propagation impairments, interference, or node mobility. Under these circumstances, the operation of an end-to-end transport protocol such as TCP may be severly affected.

Hop-by-hop transport, which distributes transport control along the source–destination path, might be considered as a "justifiable performance improvement" [3] for networks with lossy links and intermittent connectivity. In 1976, Gitman [2] compared end-to-end vs. hop-by-hop loss recovery in a scenario corresponding to an early wireless network. He found that hop-by-hop acknowledgment and retransmission leads to lower delay and higher channel utilization if there are many hops or the channel is lossy.

In this paper, we re-visit this fundamental design choice in the context of networks with considerable packet loss. Specifically, we propose a simple model of a multi-hop connection over lossy links. With this connection model, we analyze end-to-end vs. hop-by-hop retransmission in terms of delivery probability and the total number of link-level transmissions expended for the end-to-end transmission of a packet. In contrast to [2], we limit the number of retransmissions and evaluate the delivery ratio that is achieved at the expense of a certain number of link-level transmsissions. The number of link-level transmissions is a useful performance metric because it relates to both network throughput and energy consumption.

1. TRANSPORT LAYER MODEL

Our model is based on the following assumptions. A source node sends a packet to a destination node over several intermediate nodes. The nodes are connected by lossy links that have a given delivery probability. We assume un-correlated loss processes, which may represent two different things: (i) given that retransmissions occur immediately after a loss, un-correlated loss may capture link layer errors; (ii) on a large time scale, un-correlated loss may represent periods of disruption, which can be assumed to be un-correlated if a lot of time elapses between them. We consider two loss recovery schemes, namely, end-to-end and hop-by-hop transport. In the end-to-end scheme, loss recovery is a task of the source node. The source node retransmits a packet if it is lost at an intermediate hop. In the hop-by-hop scheme, loss recovery is implemented locally, i.e., each intermediate node is responsible to ensure that the packet is received by the next node. The number of transmission attempts is limited by a parameter L. In the end-to-end case, L refers to end-to-end transmission attempts from the source, and in the hop-by-hop scheme, L refers to hop-by-hop transmission attempts at the individual hops. In our derivation, we focus on the expected number of link transmission attempts. In the evaluation, we then compare both transport schemes based on this metric for values of L that result in equal end-to-end delivery ratios. Note that more complicated transport schemes can be thought of, such as those combining aspects of both end-to-end and hop-by-hop recovery mechanisms. We do not consider these schemes here for lack of space though these can also be analyzed using our approach. We use the following set of definitions:

N: number of hops
L: maximum number of transmissions allowed
p: link delivery probability
P_S: probability of successful end-to-end transmission
P_F: probability of failed end-to-end transmission
M: number of link-level transmissions

1.1 End-to-End Transport

We first determine the delivery probability over N hops with at most L end-to-end transmission attempts, denoted by P_S^{ete}. We then derive $E(M)$, the expected number of transmissions expended on the delivery of a single packet. To this end, we will first determine the expected number of transmissions given that there are Z end-to-end transmission attempts, $E(M|Z=z)$, for $z \in [1, L]$. Suppose P_l^{ete} be the probability that a packet is successfully transmitted with a maximum transmission limit of l attempts. Then $P_S^{ete} = P_L^{ete}$. Also, P_1^{ete} is the probability of success with just one end-to-end transmission attempt and equals p^N. So it follows that

$$\begin{aligned} P_S^{ete} &= 1 - (1 - P_1^{ete})^L \\ &= 1 - (1 - p^N)^L. \end{aligned}$$

In order to find $E(M|Z=z)$, we first derive $P(Z=z)$:

$$P(Z=z) = \begin{cases} (1-p^N)^{z-1} p^N & 1 \le z < L \\ (1-p^N)^{L-1} & z = L. \end{cases}$$

If we define U_i as the number of link-level transmissions in the i^{th} end-to-end transmission attempt, we can make the following observations: The expected number of link-level transmissions in all unsuccessful attempts is equal; and the number of transmissions in the successful case is equal to the number of hops of the path: $E(U_z|Z=z) = N$. We first derive $P(U|Z=z)$ for the four cases $z < Z < L, z = Z < L, z < Z = L$, and $z = Z = L$, which then allows us to express $E(U|Z=z)$ as follows:

$$E(M|Z=z) \quad = \quad E(U_1 + U_2 + \ldots + U_z)|Z=z.$$

Finally, we can write

$$E(M) \quad = \quad \sum_{z=1}^{L} P(Z=z)E(M|Z=z).$$

1.2 Hop-by-Hop Transport

We take a similar approach as in the end-to-end case, only now we condition on the number of hops traversed in a single end-to-end transmission attempt. For, if say the second hop fails then there will be no transmissions for subsequent hops. The probability of successful transmission is the probability that each hop is successful. Since the latter event happens with probability $1-(1-p)^L$, we have

$$P_S^{hbh} \quad = \quad \left(1-(1-p)^L\right)^N.$$

As before, M is the number of link-level transmissions expended per packet, and we want to find $E(M)$. We first derive the number of hops over which the packet is transmitted during one end-to-end transmission attempt, denoted by the random variable W. Let H_L be the probability that a packet is successfully relayed over one hop with a maximum of L attempts and $H_F = 1-H_L$. We have $H_L = 1-(1-p)^L$, and

$$P(W=w) \quad = \quad \begin{cases} H_L^{w-1} H_F & 1 \leq w < N \\ H_L^{w-1} & w=N. \end{cases}$$

We can now find $E(M)$ by conditioning on W. For this, we need to find $E(M|W=w)$. We can write $W=U_1+\ldots+U_W$ where U_i is the number of transmissions at the i^{th} hop. The rest of the derivation is similar to the end-to-end case and again involves distinguishing four cases depending on the relation between w, W, and L.

2. EVALUATION

We evaluate the delivery ratio P_S and the expected number of transmissions $E(M)$ for both end-to-end and hop-by-hop transport. In Fig. 1, we plot P_S against $E(M)$ with a link delivery probability of $p=0.99$ and $p=0.5$, respectively. The number of hops is $N=5$ for all experiments. The plot is generated by evaluating both metrics with $L=1\ldots300$. Please note that we consider L to be an auxiliary parameter in our derivation and we do not compare the two schemes for a given setting of L. In Fig. 1(a), both schemes achieve a delivery ratio of 0.95 with a single attempt ($L=1$) and 4.9 link-level transmissions. With a limit of two attempts, the hop-by-hop scheme reaches a ratio of 0.996 and expends 5.1 transmissions while end-to-end uses 5.15 transmissions for a ratio of 0.998. With both retransmission schemes, the delivery ratio approaches 1 for increasing values of L at the price of a marginally higher number of transmissions.

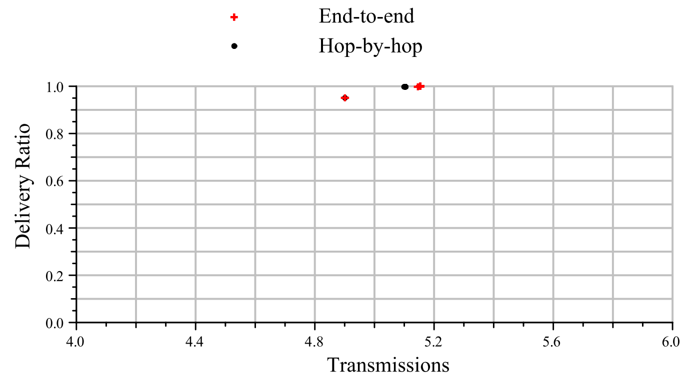

(a) p=0.99

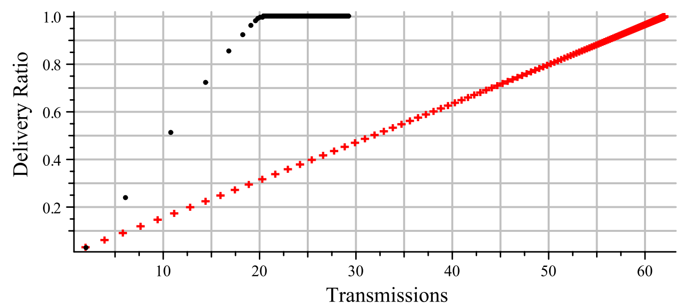

(b) p=0.5

Figure 1: Delivery ratio vs. number of transmissions

However, in Fig. 1(b), the difference between end-to-end and hop-by-hop transport is much more pronounced. With a link delivery probability of $p=0.5$, the hop-by-hop scheme approaches an end-to-end delivery ratio of 0.995 with $L=10$ and expends around 20 transmissions, corresponding to four transmissions per hop. The end-to-end scheme reaches only a ratio of 0.27 at $L=10$; for a ratio of 0.995, a setting of $L=170$ is necessary and over three times as many transmissions compared to hop-by-hop are required. Apparently, hop-by-hop is much more effective at high loss rates as it can recover from packet losses at the hop where they occur instead of retransmitting from the source.

3. CONCLUSION AND FUTURE WORK

Based on a simple model, we have investigated the efficiency in terms of link-level transmissions vs. delivery ratio of hop-by-hop against end-to-end transport. We have found that a link loss rate of 50% heavily impairs the performance of end-to-end transport even over as few as 5 hops. As a next step, we plan to introduce correlated loss processes in order to more closely model intermittently-connected networks.

4. REFERENCES

[1] D. D. Clark. The Design Philosophy of the DARPA Internet Protocols. In *SIGCOMM '88*, August 1988.

[2] I. Gitman. Comparison of hop-by-hop and end-to-end acknowledgment schemes in computer communication networks. *IEEE Transactions on Communications*, 24(11):1258–1262, 1976.

[3] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-To-End Arguments in System Design. In *ACM Transactions on Computer Systems*, November 1984.